Draft version December 20, 2016
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1612.05242v2 [astro-ph.IM] 19 Dec 2016

PAN-STARRS PHOTOMETRIC AND ASTROMETRIC CALIBRATION

Eugene. A. Magnier,[1] Edward. F. Schlafly,[2,3] Douglas P. Finkbeiner,[4,5] J. L. Tonry,[1] B. Goldman,[6] S. Röser,[7] E. Schilbach,[7] K. C. Chambers,[1] H. A. Flewelling,[1] M. E. Huber,[1] P. A. Price,[8] W. E. Sweeney,[1] C. Z. Waters,[1] L. Denneau,[1] P. Draper,[9] K. W. Hodapp,[1] R. Jedicke,[1] N. Kaiser,[1] R.-P. Kudritzki,[1] N. Metcalfe,[9] C. W. Stubbs,[10] R. J. Wainscoat[1]

Draft version December 20, 2016

ABSTRACT

The Pan-STARRS 1 3π survey has produced photometry and astrometry covering the ∼30,000 square degrees $\delta > -30°$. This article describes the photometric and astrometric calibration of this survey.

Subject headings: Surveys:Pan-STARRS 1

1. INTRODUCTION

This is the fifth in a series of seven papers describing the Pan-STARRS1 Surveys, the data reduction techiques and the resulting data products. This paper (Paper V) describes the final calibration process, and the resulting photometric and astrometric quality.

Chambers et al. (2017, Paper I) provides an overview of the Pan-STARRS System, the design and execution of the Surveys, the resulting image and catalog data products, a discussion of the overall data quality and basic characteristics, and a brief summary of important results.

Magnier et al. (2017, Paper II) describes how the various data processing stages are organised and implemented in the Imaging Processing Pipeline (IPP), including details of the the processing database which is a critical element in the IPP infrastructure .

Waters et al. (2016, Paper III) describes the details of the pixel processing algorithms, including detrending, warping, and adding (to create stacked images) and subtracting (to create difference images) and resulting image products and their properties.

Magnier et al. (2016a, Paper IV) describes the details of the source detection and photometry, including point-spread-function and extended source fitting models, and the techniques for "forced" photometry measurements.

Flewelling et al. (2016, Paper VI) describes the details of the resulting catalog data and its organization in the Pan-STARRS database. Huber et al. (2017, Paper VII) describes the Medium Deep Survey in detail, including the unique issues and data products specific to that survey. The Medium Deep Survey is not part of Data Release 1. (DR1)

The Pan-STARRS1 filters and photometric system have already been described in detail in Tonry et al. (2012).

Note: These papers are being placed on arXiv.org to provide crucial support information at the time of the public release of Data Release 1 (DR1). We expect the arXiv versions to be updated prior to submission to the Astrophysical Journal in January 2017. Feedback and suggestions for additional information from early users of the data products are welcome during the submission and refereeing process.

2. PAN-STARRS 1

From May 2010 through March 2014, the Pan-STARRS Science Consortium used the 1.8m Pan-STARRS 1 telescope to perform a set of wide-field science surveys. These surveys are designed to address a range of science goals included the search for hazardous asteroids, the study of the formation and architecture of the Milky Way galaxy, and the search for Type Ia supernovae to measure the history of the expansion of the universe.

The wide-field Pan-STARRS 1 telescope consists of a 1.8 meter diameter $f/4.4$ primary mirror with an 0.9 m secondary, producing a 3.3 degree field of view (Hodapp et al. 2004). The optical design yields low distortion and minimal vignetting even at the edges of the illuminated region. The optics, in combination with the natural seeing, result in generally good image quality: the median image quality for the 3π survey is FWHM = (1.31, 1.19, 1.11, 1.07, 1.02) arcseconds for ($g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$,$z_{\rm P1}$,$y_{\rm P1}$), with a floor of ∼ 0.7 arcseconds. The Pan-STARRS 1 camera (Tonry & Onaka 2009) is a mosaic of 60 edge-abutted 4800×4800 pixel back-illuminated CCID58 Orthogonal Transfer Arrays manufactured by Lincoln Laboratory (Tonry et al. 2006, 2008). The CCDs have 10 μm pixels subtending 0.258 arcsec and are 70μm thick. The detectors are read out using a StarGrasp CCD controller, with a readout time of 7 seconds for a full unbinned image (Onaka et al. 2008). The active, usable pixels cover ∼ 80% of the FOV.

Nightly observations are conducted remotely from the

[1] Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu HI 96822
[2] Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA
[3] Hubble Fellow
[4] Institute for Theory and Computation, Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, MS-51, Cambridge, MA 02138 USA
[5] Department of Physics, Harvard University, Cambridge, MA 02138 USA
[6] Max Planck Institute for Astronomy, Königstuhl 17, D-69117 Heidelberg, Germany
[7] Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg, Möchhofstrasse 12-14, D-69120 Heidelberg, Germany
[8] Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA
[9] Department of Physics, Durham University, South Road, Durham DH1 3LE, UK
[10] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138

Advanced Technology Research Center in Kula, the main facility of the University of Hawaii's Institute for Astronomy operations on Maui. During the Pan-STARRS 1 Science Survey, images obtained by the Pan-STARRS 1 system were stored first on computers at the summit, then copied with low latency via internet to the dedicated data analysis cluster located at the Maui High Performance Computer Center in Kihei, Maui.

Images obtained by Pan-STARRS 1 are automatically processed in real time by the Pan-STARRS 1 Image Processing Pipeline (IPP, Magnier et al. 2016a). Real-time analysis goals are aimed at feeding the discovery pipelines of the asteroid search and supernova search teams. The data obtained for the Pan-STARRS 1 Science Survey has also been used in three additional complete re-processing of the data: Processing Versions 1, 2, and 3 (PV1, PV2, and PV3). The real-time processing of the data is considered "PV0". Except as otherwise noted, the PV3 analysis of the data is used for the purpose of this article.

The data processing steps are described in detail by Waters et al. (2016) and Magnier et al. (2016a,b). In summary, individual images are detrended: non-linearity and bias corrections are applied, a dark current model is subtracted and flat-field corrections are applied. The $y_{\rm P1}$-band images are also corrected for fringing: a master fringe pattern is scaled to match the observed fringing and subtracted. Mask and variance image arrays are generated with the detrend analysis and carried forward at each stage of the IPP processing. Source detection and photometry are performed for each chip independently. As discussed below, preliminary astrometric and photometric calibrations are performed for all chips in a single exposure in a single analysis.

Chip images are geometrically transformed based on the astrometric solution into a set of pre-defined pixel grids covering the sky, called skycells. These transformed images are called the warp images. Sets of warps for a given part of the sky and in the same filter may be added together to generate deeper 'stack' images. PSF-matched difference images are generated from combinations of warps and stacks; the details of the difference images and their calibration are outside of the scope of this article.

Astronomical objects are detected and characterized in the stacks images. The details of the analysis of the sources in the stack images are discussed in Magnier et al. (2016b), but in brief these include PSF photometry, along with a range of measurements driven by the goals of understanding the galaxies in the images. Because of the significant mask fraction of the GPC1 focal plane, and the varying image quality both within and between exposures, the effective PSF of the PS1 stack images is highly variable. The PSF varies significantly on scales as small as a few to tens of pixels, making accurate PSF modelling essentially infeasible. The PSF photometry of sources in the stack images is thus degraded significantly compared to the quality of the photometry measured for the individual chip images.

To recover most of the photometric quality of the individual chip images, while also exploiting the depth afforded by the stacks, the PV3 analysis make use of forced photometry on the individual warp images. PSF photometry is measured on the warp images for all sources which are detected in the stack images images. The positions determined in the stack images are used in the warp images, but the PSF model is determined for each warp independently based on brighter stars in the warp image. The only free parameter for each object is the flux, which may be insignificant or even negative for sources which are near the faint limit of the stack detections. When the fluxes from the individual warp images are averaged, a reliable measurement of the faint source flux is determined. The details of this analysis are described in detail in Magnier et al Magnier et al. (2016b).

In this article, we discuss the photometric calibration of the individual exposures, the stacks, and the warp imags. We also discuss the astrometric calibration of the individual exposures and the stack images.

3. ASTROMETRIC MODELS

Three somewhat distinct astrometric models are employed within the IPP at different stages. The simplest model is defined independently for each chip: a simple TAN projection as described by Calabretta & Greisen (2002) is used to relate sky coordinates to a cartesian tangent-plane coordinate system. A pair of low-order polynomials are used to relate the chip pixel coordinates to this tangent-plane coordinate system. The transforming polynomials are of the form:

$$P = \sum_{i,j} C^P_{i,j} X^i_{\rm chip} Y^j_{\rm chip} \quad (1)$$

$$Q = \sum_{i,j} C^Q_{i,j} X^i_{\rm chip} Y^j_{\rm chip} \quad (2)$$

where P, Q are the tangent plane coordinates, $X_{\rm chip}, Y_{\rm chip}$ are the coordinates on the 60 GPC1 chips, and $C^P_{i,j}, C^Q_{i,j}$ are the polynomial coefficients for each order. In the `psastro` analysis, $i + j <= N_{\rm order}$ where the order of the fit, $N_{\rm order}$, may be 1 to 3, under the restriction that sufficient stars are needed to constrain the order.

A second form of astrometry model which yields somewhat higher accuracy consists of a set of connected solutions for all chips in a single exposure. This model also uses a TAN projection to relate the sky coordinates to a locally cartesian tangent plane coordinate system. A set of polynomials is then used to relate the tangent plane coordinates to a 'focal plane' coordinate system, L, M:

$$P = \sum_{i,j} C^P_{i,j} L^i M^j \quad (3)$$

$$Q = \sum_{i,j} C^Q_{i,j} L^i M^j \quad (4)$$

This set of polynomial accounts for effects such as optical distortion in the camera and distortions due to changing atmospheric refraction across the field of the camera. Since these effects are smooth across the field of the camera, a single pair of polynomials can be used for each exposure. Like in the chip analysis about, the `psastro` code restricts the exponents with the rule $i+j <= N_{\rm order}$ where the order of the fit, $N_{\rm order}$, may be 1 to 3, under the restriction that sufficient stars are needed to constrain the order For each chip, a second set of polynomials describes the transformation from the chip coordinate

systems to the focal coordinate system:

$$L = \sum_{i,j} C^L_{i,j} X^i_{\rm chip} Y^j_{\rm chip} \quad (5)$$

$$M = \sum_{i,j} C^M_{i,j} X^i_{\rm chip} Y^j_{\rm chip} \quad (6)$$

A third form of the astrometry model is used in the context of the calibration determined within the DVO database system. We retain the two levels of transformations (chip → focal plane → tangent plane), but the relationship between the chip and focal plane is represented with only the linear terms in the polynomial, supplemented by a course grid of displacements, $\delta L, \delta M$ sampled across the coordinate range of the chip. This displacement grid may have a resolution of up to 6×6 samples across the chip. The displacement for a specific chip coordinate value is determined via bilinear interpolation between the nearest sample points. Thus, the chip to focal-plane transformation may be written as:

$$L = C^L_{0,0} + C^L_{1,0} X_{\rm chip} + C^L_{0,1} Y_{\rm chip} + \delta L(X_{\rm chip}, Y_{\rm chip}) \quad (7)$$

$$M = C^M_{0,0} + C^M_{1,0} X_{\rm chip} + C^M_{0,1} Y_{\rm chip} + \delta M(X_{\rm chip}, Y_{\rm chip}) \quad (8)$$

WCS Keywords When this polynomial representation is written to the output files, a set of WCS keywords are used to define the astrometric transformation elements. It is necessary to transform the simply polynomials above into an alternate form:

$$P = \sum_{i,j} C^P_{i,j} (X_{\rm chip} - X_0)^i (Y_{\rm chip} - Y_0)^j \quad (9)$$

$$Q = \sum_{i,j} C^Q_{i,j} (X_{\rm chip} - X_0)^i (Y_{\rm chip} - Y_0)^j \quad (10)$$

4. REAL-TIME CALIBRATION

4.1. *Overview*

As images are processed by the data analysis system, every exposure is calibrated individually with respect to a photometric and astrometric database. The goal of this calibration step is to generate a preliminary astrometric calibration, to be used by the warping analysis to determine the geometric transformation of the pixels, and preliminary photometric transformation, to be used by the stacking analysis to ensure the warps are combined using consistent flux units.

The program used for the real-time calibration, `psastro`, loads the measurements of the chip detections from their individual `cmf`-format files. It uses the header information populated at the telescope to determine an initial astrometric calibration guess based on the position of the telescope boresite right ascension, declination and position angle as reported by the telescope & camera subsystems. Using the initial guess, `psastro` loads astrometric and photometric data from the reference database.

4.2. *Reference Catalogs*

During the course of the PS1SC Survey, several reference databases have been used. For the first 20 months of the survey, `psastro` used a reference catalog with synthetic PS1 $g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$,$z_{\rm P1}$,$y_{\rm P1}$ photometry generated by the Pan-STARRS IPP team based on based combined photometry from Tycho (B, V), USNO (red, blue, IR), and 2MASS J, H, K. The astrometry in the database was from 2MASS. After 2012 May, a reference catalog generated from internal re-calibration of the PV0 analysis of PS1 photometry and astrometry was used for the reference catalog.

Coordinates and calibrated magnitudes of stars from the reference database are loaded by `pasastro`. A model for the positions of the 60 chips in the focal plane is used to determine the expected astrometry for each chip based on the boresite coordinates and position angle reported by the header. Reference stars are selected from the full field of view of the GPC1 camera, padded by an additional 25% to ensure a match can be determined even in the presence of substantial errors in the boresite coordinates. It is important to choose an appropriate set of reference stars: if too few are selected, the chance of finding a match between the reference and observed stars is diminished. In addition, since stars are loaded in brightness order, a selection which is too small is likely to contain only stars which are saturated in the GPC1 images. On the other hand, if too many reference stars are chosen, there is a higher chance of a false-positive match, especially as many of the reference stars may not be detected in the GPC1 image. The seletion of the reference stars includes a limit on the brightest and fainted magnitude of the stars selected.

The astrometric analysis is necessarily performed first; after the astrometry is determined, an automatic byproduct is a reliable match between reference and observed stars, allowing a comparison of the magnitudes to determine the photometric calibration.

The astrometric calibration is performed in two major stages: first, the chips are fitted independently with independent models for each chip. This fit is sufficient to ensure a reliable match between reference stars and observed sources in the image. Next, the set of chip calibrations are used to define the transformation between the focal plane coordinate system and the tangent plane coordinate system. The chip-to-focal plane transformations are then determined under the single common focal plane to tangent plane transformation.

4.3. *Cross-Correlation Search*

The first step of the analysis is to attempt to find the match between the reference stars and the detected objects. `psastro` uses 2D cross correlation to search for the match. The guess astrometry calibration is used to define a predicted set of $X^{\rm ref}_{\rm chip}, Y^{\rm ref}_{\rm chip}$ values for the reference catalog stars. For all possible pairs between the two lists, the values of

$$\Delta X = X^{\rm ref}_{\rm chip} - X^{\rm obs}_{\rm chip} \quad (11)$$

$$\Delta Y = Y^{\rm ref}_{\rm chip} - Y^{\rm obs}_{\rm chip} \quad (12)$$

are generated. The collection of $\Delta X, \Delta Y$ values are collected in a 2D histogram with sampling of 50 pixels and the peak pixel is identified. If the astrometry guess were perfect, this peak pixel would be expected to lie at (0,0) and contain all of the matched stars. However, the astrometric guess may be wrong in several ways. An error in the constant term above, $C^P_{0,0}, C^Q_{0,0}$ shifts the peak to another pixel, from which $C^P_{0,0}, C^Q_{0,0}$ can easily be deter-

mined. An error in the plate scale or a rotation will smear out the peak pixel potentially across many pixels in the 2D histogram.

To find a good match in the face of plate scale and rotation errors, the cross correlation analysis above is performed for a series of trials in which the scale and rotation are perturbed from the nominal value by a small amount. For each trial, the peak pixel is found and a figure of merit is measured. The figure of merit is defined as $\frac{\sigma_x^2+\sigma_y^2}{N_p^4}$ where $\sigma_{x,y}^2$ are the second moment of $\Delta X, Y$ for the star pairs associated with the peak pixel, and N_p is the number of star pairs in the peak. This figure of merit is thus most sensitive to a narrow distribution with many matched pairs. For the PS1 exposures, rotation offsets of (-1.0, -0.5, 0.0, 0.5, 1.0) degrees, and plate scales of (+1%, 0, -1%) of the nominal plate scale are tested. The best match among these 15 cross-correlation tests is selected and used to generate a better astrometry guess for the chip.

4.4. *Chip Polynomial Fits*

The astrometry solution from the cross correlation step above is again used to selected matches between the reference stars and observed stars in the image. The matching radius starts off quite large, and a series of fits is performed to generate the transformation between chip and tangent plane coordinates. Three clipping iterations are performed, with outliers $> 3\sigma$ rejected on each pass, where here σ is determined from the distribution of the residuals in each dimension (X,Y) independently. After each fit cycle, the matches are redetermined using a smaller radius and the fit re-tried.

4.5. *Mosaic Astrometry Polynomial Fits*

The astrometry solutions from the independent chip fits are used to generate a single model for the camera-wide distortion terms. The goal is to determine the two stage fit (chip $\rightarrow$ focal plane $\rightarrow$ tangent plane). There are a number of degenerate terms between these two levels of transformation, most obviously between the parameters which define the constant offset from chip to focal plane ($C_{0,0}^{L,M}$) and those which define the offset from focal plane to tangent plane ($C_{0,0}^{P,Q}$). We limit ($C_{0,0}^{P,Q}$) to be 0,0 to remove this degeneracy.

The initial fit of the astrometry for each chip follows the distortion introduced by the camera: the apparent plate scale for each chip is the combination of the plate scale at the optical axis of the camera, modified by the local average distortion. To isolate the effect of distortion, we choose a single common plate scale for the set of chips and re-define the chip $\rightarrow$ sky calibrations as a set of chip $\rightarrow$ focal plane transformation using that common pixel scale. We can now compare the observed focal plane coordinates, derived from the chip coordinates, and the tangent plane coordiantes, derived from the projection of the reference coordinates. One caveat is that the chip reference coordinates are also degenerate with the fitted distortion. In order to avoid being sensitive to the exact positions of the chips at this stage, we measure the local gradient between the focal plane and tangent plane coordinate systems. We then fit the gradient with a polynomial of order 1 less than the polynomial desired for the distortion fit. The coefficients of the gradient fit are then used to determine the coefficients for the polynomials representing the distortion.

Once the common distortion coming from the optics and atmosphere have been modeled, `psastro` determines polynomial transformations from the 60 chips to the focal plane coordinate system. In this stage, 5 iterations of the chip fits are performed. Before each iteration, the reference stars and detected objects are matched using the current best set of transformations. These fits start with low order (1) and large matching radius. As the iterations proceed, the radius is reduced and the order is allowed to increaes, up to 3rd order for the final iterations.

4.6. *Real-time Photometric Calibration*

After the astrometric calibration has finished, the photometric calibration is performed by `psastro`. When the reference stars are loaded, the apparent magnitude in the filter of interest is also loaded. Stars for which the reference magnitude is brighter than ($g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$,$z_{\rm P1}$,$y_{\rm P1}$) = (19, 19, 18.5, 18.5, 17.5) are used to determine the zero points by comparison with the instrumental magnitudes. For the PV3 analysis, an outlier-rejecting median is used to measure the zero point. For early versions of the analysis, when the reference catalog used synthetic magnitudes, it was necessary to search for the blue edge of the distribution: the synthetic magnitude poorly predicted the magnitudes of stars in the presence of significant extinction or for the very red stars, making the blue edge somewhat more reliable. Note that we do not include an airmass correction in this zero point analysis: the airmass correction is folded into the observed zero point. The zero point may be measured separately for each chip or as a single value for the entire exposure; the latter option was used for the PV3 analysis.

4.7. *Real-time outputs*

The calibrations determined by `psastro` as saved as part of the header information in the output FITS tables. A single multi-extension FITS table is written using the `smf` format. In these files, the measurements from each chip are written as a separate FITS table. A second FITS extension for each chip is used to store the header information from the original chip image. The original chip header is modified so that the extension corresponds to an image with no pixels data: `NAXIS` is set to 0, even though `NAXIS1` and `NAXIS2` are retained with the original dimensions of the chip. A pixel-less primary header unit (PHU) is generated with a summary of some of the important and common chip-level keywords (e.g., `DATE-OBS`). The astrometric transformation information for each chip is saved in the corresponding header using standard (and some non-standard) WCS keywords. For the two-level astrometric model, the PHU header carries the astrometric transformation related to the projection and the camera-wide distortions. Photometric calibrations are written as a set of keywords to individual chip headers, and if the calibration is performed at the exposure-level, to the PHU. The photometry calibration keywords are:

- `ZPT_REF` : the nominal zero point for this filter

- `ZPT_OBS` : the measured zero point for this chip / exposure
- `ZPT_ERR` : the measured error on `ZPT_OBS`
- `ZPT_NREF` : the number of stars used to measure `ZPT_OBS`
- `ZPT_MIN` : minimum reference magnitude included in analysis
- `ZPT_MAX` : maximum reference magnitude included in analysis

The keyword `ZPT_OBS` is used to set the initial zero point when the data from the exposure are loaded into the DVO database.

5. PV3 DVO MASTER DATABASE

Data from the GPC1 chip images, the stack images, and the warp images are loaded into DVO using the real-time analysis astrometric calibration to guide the association of detections into objects. After the full PV3 DVO database was constructed, including all of the chip, stack, and warp detections, several external catalogs were merged into the database. First, the complete 2MASS PSC was loaded into a stand-alone DVO database, which was then merged into the PV3 master database. Next the DVO database of synthetic photometry in the PS1 bands (see Section **??**) was merged in. Next, the full Tycho database was added, followed by the AllWISE database. After the Gaia release in August 2016 (Gaia Collaboration et al. 2016), we generated a DVO database of the Gaia positional and photometric information and merged that into the master DVO database.

6. PHOTOMETRY CALIBRATION

6.1. *Ubercal Analysis*

The photometric calibration of the DVO database starts with the "ubercal" analysis technique as described by Schlafly et al. (2012). This analysis is performed by the group at Harvard, loading data from the `smf` files into their instance of the Large Scale Database (LSD, Juric 2011), a system similar to DVO used to manage the detections and determine the calibrations.

Photometric nights are selected and all other exposures are ignored. Each night is allowed to have a single fitted zero point and a single fitted value for the airmass extinction coefficient per filter. The zero points and extinction terms are determined as a least squares minimization process using the repeated measurements of the same stars from different nights to tie nights together. Flat-field corrections are also determined as part of the minimization process. In the original (PV1) ubercal analysis, Schlafly et al. (2012) determined flat-field corrections for 2×2 sub-regions of each chip in the camera and four distinct time periods ("seasons"). Later analysis (PV2) used an 8×8 grid of flat-field corrections to good effect.

The ubercal analysis was re-run for PV3 by the Harvard group. For the PV3 analysis, under the pressure of time to complete the analysis, we chose to use only a 2×2 grid per chip as part of the ubercal fit and to leave higher frequency structures to the later analysis. A 5th flat-field season consisting of nearly the last 2 years of data was also included for PV3. In retrospect, as we show below, the data from the latter part of the survey would probably benefit from additional flat-field seasons.

By excluding non-photometric data and only fitting 2 parameters for each night, the Ubercal solution is robust and rigid. It is not subject to unexpected drift or sensitivity of the solution to the vagaries of the data set. The Ubercal analysis is also especially aided by the inclusion of multiple Medium Deep field observations every night, helping to tie down overall variations of the system throughput and acting as internal standard star fields. The resulting photometric system is shown by Schlafly et al. (2012) to have reliability across the survey region at the level of (8.0, 7.0, 9.0, 10.7, 12.4) millimags in ($g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$,$z_{\rm P1}$,$y_{\rm P1}$). As we discuss below, this conclusion is reinforced by our external comparison.

The overall zero point for each filter is not naturally determined by the Ubercal analysis; an external constraint on the overall photometric system is required for each filter. Schlafly et al. (2012) used photometry of the MD09 Medium Deep field to match the photometry measured by Tonry et al. (2012) on the reference photometric night of MJD 55744 (UT 02 July 2011). Scolnic et al. (2014) and Scolnic et al. (2015) have re-examined the photometry of Calspec standards observed by PS1. Scolnic et al. (2014) reject 2 of the 7 stars used by Tonry et al. (2012) and add photometry of 5 additional stars. Scolnic et al. (2015) further reject measurements of Calspec standards obtained close to the center of the camera field of view where the PSF size and shape changes very rapidly. The result of this analysis modifies the over system zero points by 20 - 35 millimags compared with the system determined by Schlafly et al. (2012).

6.2. *Applying the Ubercal Zero Points : Setphot*

The ubercal analysis above results in a table of zero points for all exposures considered to be photometric, along with a set of low-resolution flat-field corrections. It is now necessary to use this information to determine zero points for the remaining exposures and to improve the resolution of the flat-field correction. This analysis is done within the IPP DVO database system.

The ubercal zero points and the flat-field correction data are loaded into the PV3 DVO database using the program `setphot`. This program converts the reported zero point and flat field values to the DVO internal representation in which the zero point of each image is split into three main components:

$$zp_{\rm total} = zp_{\rm nominal} + M_{cal} + K_{rm\lambda}(sec\zeta - 1)$$

where $zp_{\rm nominal}$ and $K_{rm\lambda}$ are static values for each filter representing respectively the nominal zero point and the slope of the trend with respect to the airmass (ζ) for each filter. These static values are listed in Table **??**. When `setphot` was run, these static zero points have been adjusted by the Calspec offsets listed in Table **??** based on the analysis of Calspec standards by Scolnic et al. (2015). These offsets bring the photometric system defined by the ubercal analysis into alignment with Scolnic et al. (2015). The value M_{cal} is the offset needed by each exposure to match the ubercal value, or to bring the non-ubercal exposures into agreement with the rest of the exposures, as discussed below. The flat-field information is encoded in a table of flat-field offsets as a

function of time, filter, and camera position. Each image which is part of the ubercal subset is marked with a bit in the field `Image.flags`: `ID_IMAGE_PHOTOM_UBERCAL = 0x00000200`

When `setphot` applies the ubercal information to the image tables, it also updates the individual measurements associated with those images. In the DVO database schema, the normalized instrumental magnitude, $M_{\rm inst} = -2.5log_{10}(DN/sec) + 25.0$ are stored for each measurement. The value of 25.0 is an arbitrary (but fixed) constant offset to place the instrumental magnitudes into approximately the correct range. Associated with each measurement are two correction magnitudes: $M_{\rm cal}$ and $M_{\rm flat}$, along with the airmass for the measurement, calculated using the altitude of the individual detection as determined from the Right Ascension, Declination, the observatory latitude, and the sidereal time. For a camera with the field of view of the PS1 GPC1, the airmass may vary significantly within the field of view, especially at low elevations. In the worst cases, at the celestial pole, the airmass range within a single exposure is XXX - XXX. The complete calibrated ('relative') magnitude is determined from the stored database values as:

$$M_{\rm rel} = M_{\rm inst} - 25.0 + zp_{\rm ref} + M_{\rm cal} + M_{\rm flat} + K_\lambda(sec\zeta - 1).$$

The calibration offsets, $M_{\rm cal}$ and $M_{\rm flat}$, represent the per-exposure zero point correction and the slowly-changing flat-field correction respectively. These two values are split so the flat-field corrections may be determined and applied independently from the time-resolved zero point variations. Note that the above corrections are applied to each of the types of measurements stored in the database, PSF, Aperture, Kron. The calibration math remains the same regardless of the kind of magnitude being measured. Also note that for the moment, this discussion should only be considered as relevant to the chip measurements. Below we discuss the implications for the stack and warp measurements.

When the ubercal zero points and flat-field data are loaded, `setphot` updates the $M_{\rm cal}$ values for all measurements which have been derived from the ubercal images. These measurements are also marked in the field `Measure.dbFlags` with the bit `ID_MEAS_PHOTOM_UBERCAL = 0x00008000`. At this stage, `setphot` also updates the values of $M_{\rm flat}$ for all GPC1 measurements in the appropriate filters.

6.3. *Relphot Analysis*

Relative photometry is used to determine the zero points of the exposures which were not included in the ubercal analysis. The relative photometry analysis has been described in the past by Magnier et al. (2013). We review that analysis here, along with specific updates for PV3.

As described above, the instrumental magnitude and the calibrated magnitude are related by arithmetic magnitude offsets which account for effects such as the instrumental variations and atmospheric attenuation.

$$M_{rel} = m_{inst} + ZP + M_{cal}$$

From the collection of measurements, we can generate an average magnitude for a single star (or other object):

$$M_{ave} = \frac{\sum_i M_{rel,i} w_i}{\sum_i w_i}$$

We find that the color difference of the different chips can be ignored, and set the value of A to 0.0. Note that we only use a single mean airmass extinction term for all exposures – the difference between the mean and the specific value for a given night is taken up as an additional element of the atmospheric attenuation.

We write a global χ^2 equation which we attempt to minimize by finding the best mean magnitudes for all objects and the best $M_{\rm cal}$ offset for each exposure:

$$\chi^2 = \sum_{i,j} (m_{inst}[i,j] + ZP + K\zeta + M_{clouds}[i] - M_{ave}[j]) w_{i,j} / \sum_{i,j} w_{i,j}$$

If everything were fitted at once and allowed to float, this system of equations would have $N_{exposures} + N_{stars} \sim 2 \times 10^5 + N \times 10^9$ unknowns. We solve the system of equations by iteration, solving first for the best set of mean magnitudes in the assumption of zero clouds, then solving for the clouds implied by the differences from these mean magnitudes. Even with 1-2 magnitudes of extinction, the offsets converge to the milli-magnitude level within 8 iterations.

Only brighter, high quality measurements are used in the relative photometry analysis of the exposure zero points. We use only the brighter objects, limiting the density to a maximum of 4000 objects per square degree (lower in areas where we have more observations). When limiting the density, we prefer objects which are brighter (but not saturated), and those with the most measurements (to ensure better coverage over the available images).

There are a few classes of outliers which we need to be careful to detect and avoid. First, any single measurement may be deviant for a number of reasons (e.g., it lands in a bad region of the detector, contamination by a diffraction spike or other optical artifact, etc). We attempt to exclude these poor measurements in advance by rejecting measurements which the photometric analysis has flagged the result as suspcious. We reject detections which are excessively masked; these include detections which are too close to other bright objects, diffraction spikes, ghost images, or the detector edges. However, these rejections do not catch all cases of bad measurements.

After the initial iterations, we also perform outlier rejections based on the consistency of the measurements. For each star, we use a two pass outlier clipping process. We first define a robust median and sigma from the inner 50% of the measurements. Measurements which are more than 5σ from this median value are rejected, and the mean & standard deviation (weighted by the inverse error) are recalculated. We then reject detections which are more than 3σ from the recalculated mean.

Suspicious stars are also exclude from the analsis. We exclude stars with reduced χ^2 values more than 20.0, or more than $2\times$ the median, whichever is larger. We also exclude stars with standard deviation (of the measurements used for the mean) greater than 0.005 mags or $2\times$ the median standard deviation, whichever is greater.

Similarly for images, we exclude those with more than 2 magnitudes of extinction or for which the deviation greater of the zero points per star are than 0.075 mags or 2× the median value, whichever is greater. These cuts are somewhat conservative to limit us to only good measurements. The images and stars rejected above are not used to calculate the system of zero points and mean magnitudes. These cuts are updated several times as the iterations proceed. After the iterations have completed, the images which have been reject are calibrated based on their overlaps with other images.

We overweight the ubercal measurements in order to tie the relative photometry system to the ubercal zero points. Ubercal images and measurements from those images are not allowed to float in the relative photometry analysis. Detections from the Ubercal images are assigned weights of 10x their default (inverse-variance) weight. The calculation of the formal error on the mean magnitudes propagates this additional weight, so that the errors on the Ubercal observations dominates where they are present.

$$\mu = \frac{\sum m_i w_i \sigma_i^{-2}}{\sum w_i \sigma_i^{-2}}$$

$$\sigma_\mu = \frac{\sum w_i^2 \sigma_i^{-2}}{(\sum w_i \sigma_i^{-2})^2}$$

The calculation of the relative photometry zero points is performed for the entire 3π data set in a single, highly parallelized analysis. As discussed above, the measurement and object data in the DVO database are distributed across a large number of computers in the IPP cluster: for PV3, 100 parallel hosts are used. These machines by design control data from a large number of unconnected small patches on the sky, with the goal of speeding queries for arbitrary chunks of the sky. As a result, this parallelization is entirely inappropriate as the basis of the relative photometry analysis. For the relative photometry calculation (and later for relative astrometry calculation), the sky is divided into a number of large, contiguous regions each bounded by lines of constant RA & DEC, 73 regions in the case of the PV3 analysis. A separate computer, called a "region host" is responsible for each of these regions: that computer is responsible for calculating the mean magnitudes of the objects which land within its region and for determining the exposure zero points for exposures for which the center of the exposure lands in the region of responsibility.

The iterations described above (calculate mean magnitudes, calculate zero points, calculate new measurements) are peformed on each of the 73 region hosts in parallel. However, between certain iteration steps, the region hosts must share some information. After mean object magnitudes are calculated, the region hosts must share the object magnitudes for the objects which are observed by exposures controlled by neighboring region hosts. After image calibrations have been determined by each region host, the image calibrations must be shared with the neighboring region hosts so measurement values associated with objects owned by a neighboring region host may be updated.

The completely work flow of the all-sky relative photometry analysis starts with an instance of the program running on a master computer. This machine loads the image database table and assigns the images to the 73 region hosts. A process is then launched on each of the region hosts which is responsible for managing the image calibration analysis on that host. These processes in turn make an intial request of the photometry information (object and measurement) from the 100 parallel DVO partition machines. In practice, the processes on the the region hosts are launched in series by the master process to avoid overloading the DVO partition machines with requests for photometry data from all region hosts at once. Once all of the photometry has been loaded, the region hosts perform their iterations, sharing the data which they need to share with their neighbors and blocking while they wait for the data they need to receive from their neighbors. The management of this stage is performed by communication between the region host. At the end of the iterations, the regions hosts write out their final image calibrations. The master machine then loads the full set of image calibrations and then applies these calibrations back to all measurements in the database, updating the mean photometry as part of this process. The calculations for this last step are performed in parallel on the DVO partition machines.

With the above software, we are able to perform the entire relphot analysis for the full 3π region at once, avoiding any possible edge effects. The region host machines have internal memory ranging from 96GB to 192GB. Regions are drawn, and the maximum allowed density was chosen, to match the memory usage to the memory available on each machine. A total of 9.8TB of RAM was available for the analysis, allowing for up to 6000 objects per square degree in the analysis.

6.3.1. *Photometric Flat-field*

For PV3, the relphot analysis was performed two times. The first analysis used only the flat-field corrections determined by the ubercal analysis, with a resolution of 2x2 flat-field values for each GPC1 chip (corresponding to ∼2400 pixels), and 5 separate flat-field 'seasons'. However, we knew from prior studies that there were significant flat-field structures on smaller scales. We used the data in DVO after the initial relphot calibration to measure the flat-field residual with much finer resolution: 124 x 124 flat-field values for each GPC1 chip (40x40 pixels per point). We then used `setphot` to apply this new flat-field correction, as well as the ubercal flat-field corrections, to the data in the database. At this point, we re-ran the entire relphot analysis to determine zero points and to set the average magnitudes.

Figure 1 shows the high-resolution photometric flat-field corrections applied to the measurements in the DVO database. These flat-fields make low-level corrections of up to ∼0.03 magnitudes. Several features of interest are apparent in these images.

First, at the center of the camera is an important structure caused by the telescope optics which we call the "tent". In this portion of the focal plane, the image quality degrades very quickly. The photometry is systematically biased because the point spread function model cannot follow the real changes in the PSF shape on these small scales. As is evident in the image, the effect is such that the flux measured using a PSF model is systematically low, as expected if the PSF model is too

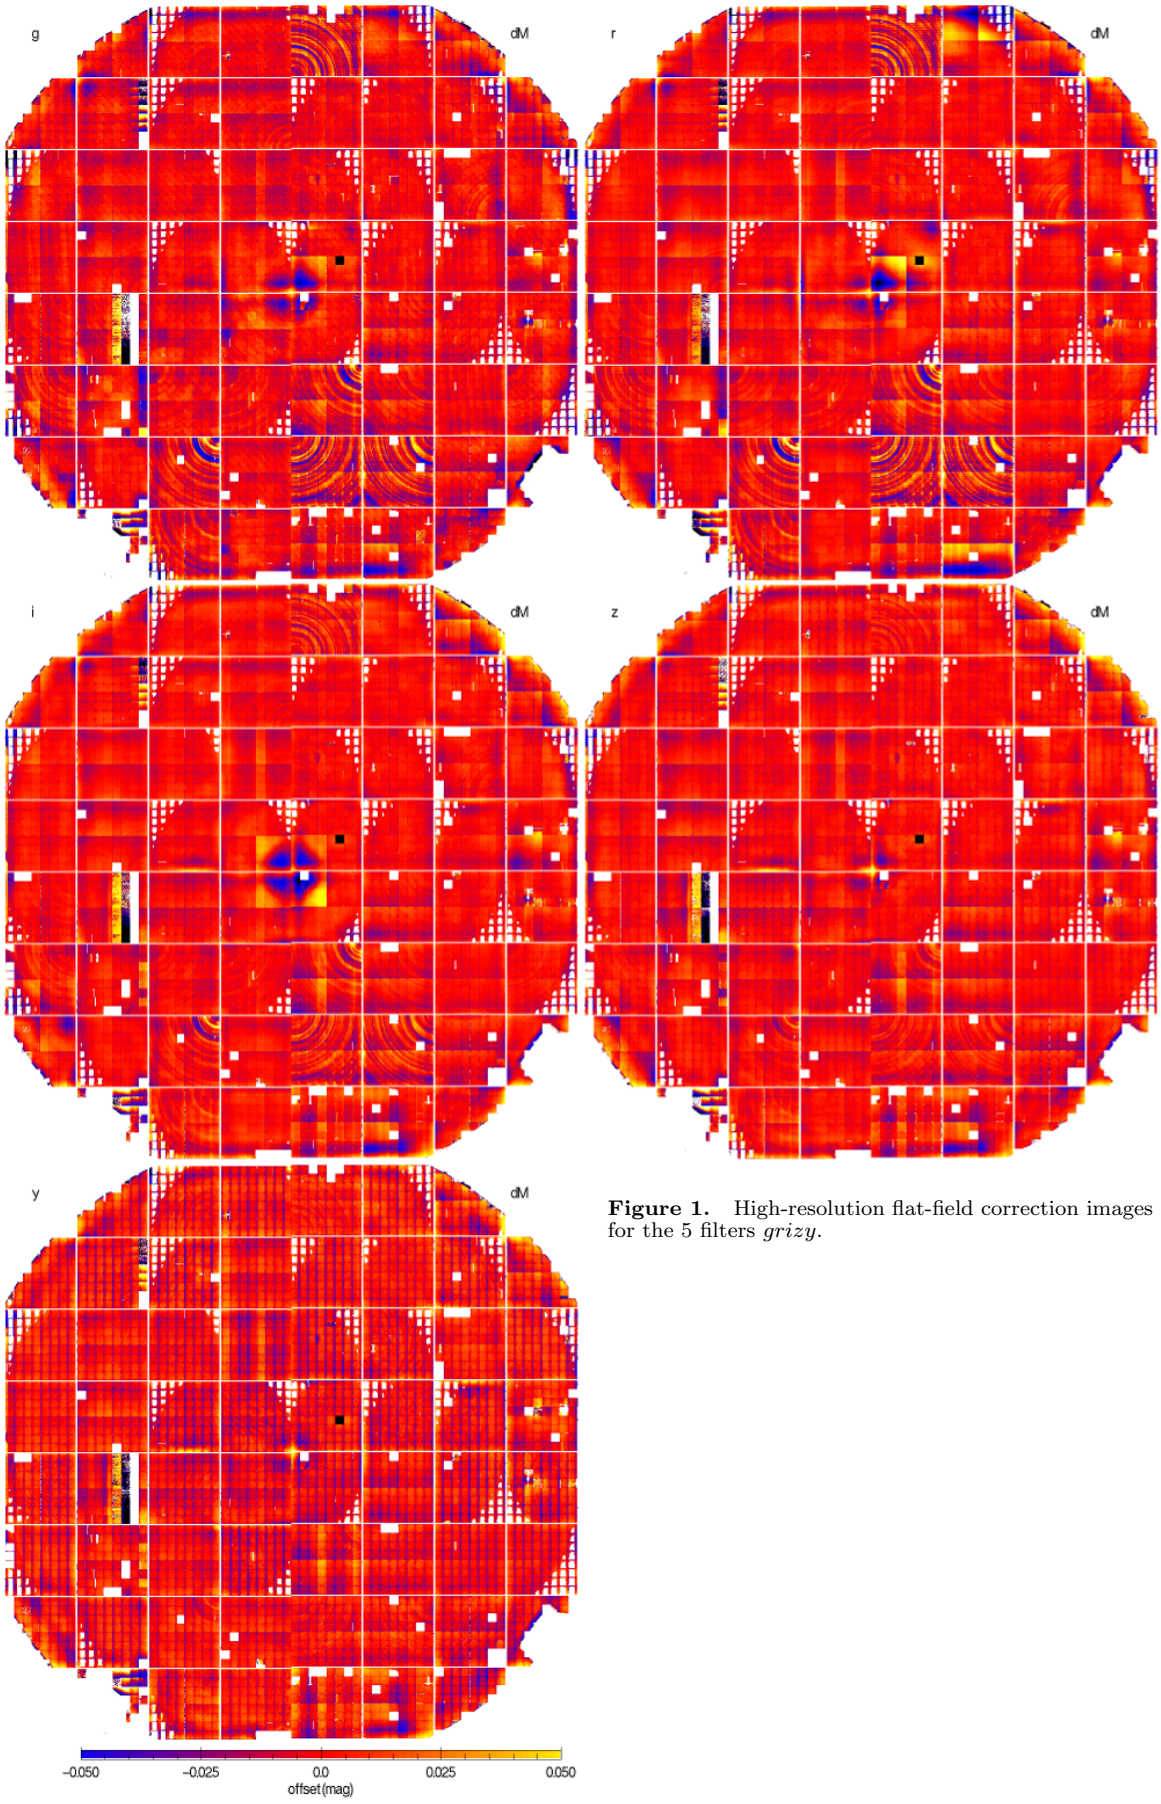


Figure 1. High-resolution flat-field correction images for the 5 filters *grizy*.

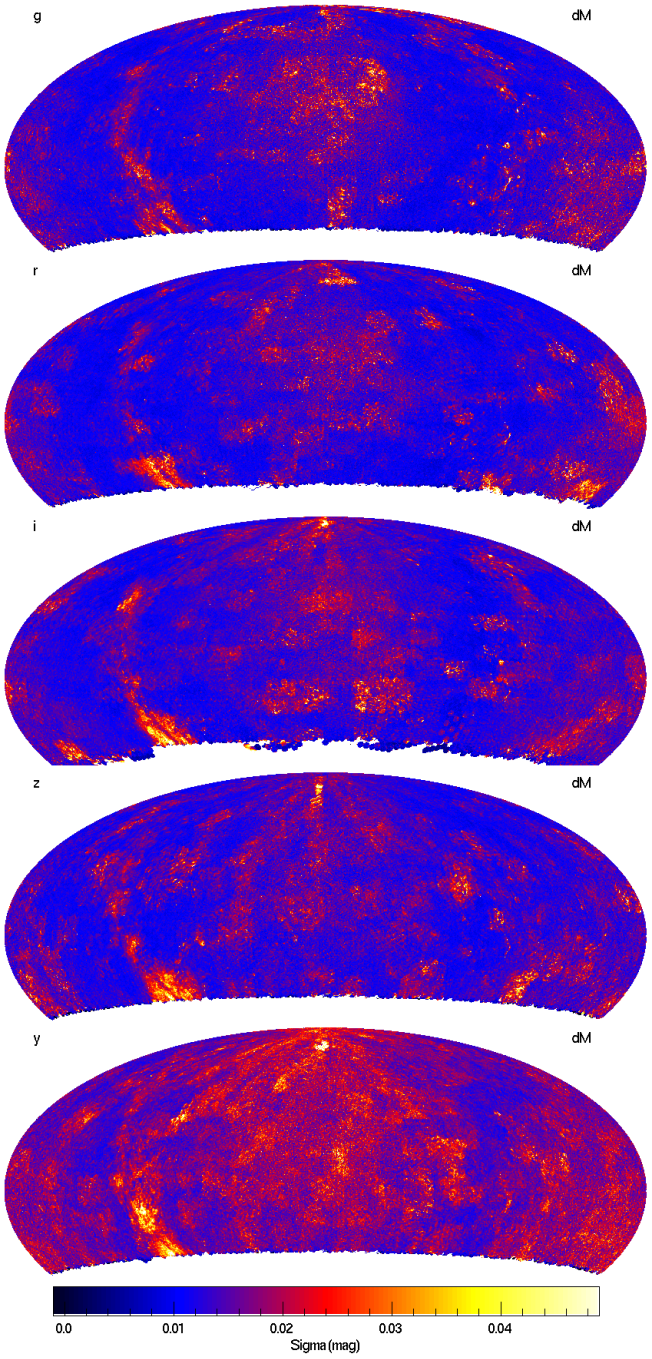


Figure 2. Consistency of photometry measurements across the sky. Each panel shows a map of the standard deviation of photometry residuals for stars in each pixel. The median value of the measure standard deviations across the sky is $(\sigma_g, \sigma_r, \sigma_i, \sigma_z, \sigma_y) = (14, 14, 15, 15, 18)$ millimags. These values reflect the typical single-measurement errors for bright stars.

small.

The square outline surrounding the "tent" is due to the 2×2 sampling per chip used for the Ubercal flat-field corrections. The imprint of the Ubercal flat-field is visible throughout this high-resolution flat-field: in regions where the underlying flat-field structure follows a smooth gradient across a chip, the Ubercal flat-field partly corrects the structure, leaving behind a saw-tooth residual. The high-resolution flat-field corrects the residual structures well.

Especially notable in the bluer filters is a pattern of quarter circles centered on the corners of the chips. These patterns are similar to the "tree rings" reported by the DES team and others (G. Berstein REF & REFS). The details of these tree rings are beyond the scope of this article, and will be explored in future work. Unlike the tree ring features discussed by these other authors, the features observed in the GPC1 photometry are not caused by an interaction of the flat-field with the effective pixel geometry. Instead, these photometric features are due to low-level changes in the PSF size which we attribute to variable charge diffusion (Magnier in prep).

Other features include some poorly responding cells (e.g., in XY14) and effects at the edges of chips, possibly where the PSF model fails to follow the changes in the PSF.

For stacks and warps, the image calibrations were determined after the relative photometry was performed on the individual chips. Each stack and each warp was tied via relative photometry to the average magnitudes from the chip photometry. In this case, no flat-field corrections were applied. For the stacks, such a correction would not be possible after the stack has been generated because multiple chip coordinates contribute to each stack pixel coordinate. For the warps, it is in principle possible to map back to the corresponding chip, but the information was not available in the DVO database, and thus it was not possible at this time to determine the flat-field correction appropriate for a given warp. This latter effect is one of several which degrade the warp photometry compared to the chip photometry at the bright end.

6.4. *Photometry Calibration Quality*

Figure 2 shows the standard devitions of the mean residual photometry for bright stars as a function of position across the sky. For each pixel in these images, we selected all objects with (14.5, 14.5, 14.5, 14.0, 13.0) $<$ (g, r, i, z, y) $<$ (17, 17, 17, 16.5, 15.5), with at least 3 measurements in i-band (to reject artifacts detected in a pair of exposures from the same night), with `PSF_QF` > 0.85 (to reject excessively-masked objects), and with $mag_{\rm PSF} - mag_{rmKron} < 0.1$ (to reject galaxies). We then generated histograms of the difference between the average magnitude and the apparent magnitude in an individual image for each filter for all stars in a given pixel in the images. From these residual histograms, we can then determine the median and the 68%-ile range to calculate a robust standard deviation. This represents the bright-end systematic error floor for a measurement from a single exposure. The standard deviations are then plotted in Figure 2.

The 5 panels in Figure 2 show several features. The Galactic bulge is clearly seen in all five filters, with the impact strongest in the reddest bands. We attribute this to the effects of crowding and contamination of the photometry by neighbors. Large-scale, roughly square features ∼10 degrees on a side in these images can be attributed to the vagaries of weather: these patches correspond to the observing chunks. These images include both photometric and non-photometric exposures. It seems plausible that the non-photometric images from relatively poor quality nights elevate the typical errors. On small scales, there are circular pat-

terns ∼3 degrees in diameter corresponding to individual exposures; these represent residual flat-fields structures not corrected by our stellar flat-fielding. The median of the standard deviations in the five filters are $(\sigma_g, \sigma_r, \sigma_i, \sigma_z, \sigma_y) = (14, 14, 15, 15, 18)$ millimagnitudes.

6.5. *Calculation of Object Photometry*

6.5.1. *Iteratively Reweighted Least Squares Fitting (1D)*

6.5.2. *Selection of Measurements*

6.5.3. *Stack Photometry*

6.5.4. *Warp Photometry*

7. ASTROMETRY CALIBRATION

Once the full PV3 dataset loaded into the master PV3 DVO database, along with supporting databases, and the photometric calibrations were performed, relative astrometry could be performed on the database to improve the overall astrometric calibration.

In many respects the relative astrometric analysis is similar to the relative photometric analysis: the repeated measurements of the same object in different images are used to determine a high quality average position for the object. The new average positions are then used to determine improved astrometric calibrations for each of the images. These improved calibrations are used to set the observed coordinates of the measurements from those images, which are in turn used to improve the average positions of the objects. The whole process is repeated for several iterations. Like the photometric analysis, the astrometric analysis is performed in a parallel fashion with the same concept that specific machines are responsible for exposures and objects which land within their regions of responsibility, defined on the basis of lines of constant RA and DEC. Between iteration steps, the astrometric calibrations are shared between the parallel machines as are the improved positions for objects controlled by one machine but detect in images controlled by another machine. Like the photometric analysis, the entire sky is processed in one pass. However, there are some important differences in the details.

7.1. *Systematic Effects*

First, the astrometric calibration has a larger number of systematic effects which must be performed. These consist of: 1) the Koppenhöffer Effect, 2) Differential Chromatic Refraction, 3) Static deviations in the camera. We discuss each of these in turn below.

7.1.1. *Koppenhöffer Effect*

The Koppenhöffer Effect was first identified in February 2011 by Johannes Koppenhöffer (MPE) as part of the effort to search for planet transists in the Stellar Transit Survey data. He noticed that the astromety of bright stars and faint stars disagreed on overlapping chips at the boundary between the STS fields. After some exploration, it was determined that the X coordinate of the brightest stars was offset from the expected location based on the faint stars for a subset of the GPC1 chips. The essence of the effect was that a large charge packet could be drawn prematurely over an intervening negative serial phase into the summing well, and this leakage was proportionately worse for brighter stars. The brighter the star, the more the charge packet was pushed ahead on the serial register. The amplitude of the effect was at most $0\farcs25$, corresponding to a shift of about one pixel. This effect was only observed in 2-phase OTA devices, with 22 / 30 of these suffering from this effect. By adjusting the summing well high voltage down from a default +7 V to +5.5V on the 2-phase devices, the effect was prevented in exposures after 2011-05-03. However, this left 101,550 exposures (27%) already contaminated by the effect.

We measured the Koppenhöffer Effect by accumulating the residual astrometry statistics for stars in the database. For each chip, we measured the mean X and Y displacements of the astrometric residuals as function of the instrumental magnitude of the star divided by the FWHM^2. We measured the trend for all chips in a number of different time ranges and found the effect to be quite stable, in the period where it was present. The effect only appeared in the serial direction. Figure **??** shows the KE trend for a typical affected chip both before and after the correction. For the PV3 dataset, we re-measured the KE trends using stars in the Galactic pole regions after an initial relative astrometry calibration pass: the Galactic pole is necessary because the real-time astrometric calibration relies largely on the fainter stars which are not affected by the KE. The trend is then stored in a form which can be applied to the database measurements.

7.1.2. *Differential Chromatic Refraction*

Differential Chromatic Refraction (DCR) affects astrometry because the reference stars used the calibrate the images are not the same color (SED) as the rest of the stars in the image. For a given star of a color different from the reference stars, as exposures are taken at higher airmass, the apparent position of the star will be biased along the parallactic angle. While it is possible to build a model for the DCR impact based on the filter response functions and atmospheric refraction, we have instead elected to use an empirical correction for the DCR present in the PV3 database. We have measured the DCR trend using the astrometric residuals of millions of stars after performing an initial relative astrometry calibration. We define a blue DCR color $(g-i)$ to be used when correcting the filters $g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$, and a red DCR color $(z-y)$ to be used when correcting the filters zy. In the process of performing the relative astrometry calibration, we record the median red and blue colors of the reference stars used to measure the astrometry calibration for each image. As we determine the astrometry parameters for each object in the database, we record the median red and blue reference star colors for all images used to determine the astrometry for a given object. For each star in the database, we know both the color of the star and the typical color of the reference stars used to calibrate the astrometry for that star.

We measure the mean deviation of the residuals in the parallactic angle direction and the direction perpendicular to the parallactic angle. For each filter, we determine the DCR trend as a function of the difference between the star color and the reference star color, using the red or blue color approriate to the particular filter, times the tangent of the zenith distance. Figure **??** shows the DCR trend for the 5 filters $g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$,$z_{\rm P1}$,$y_{\rm P1}$, as well as

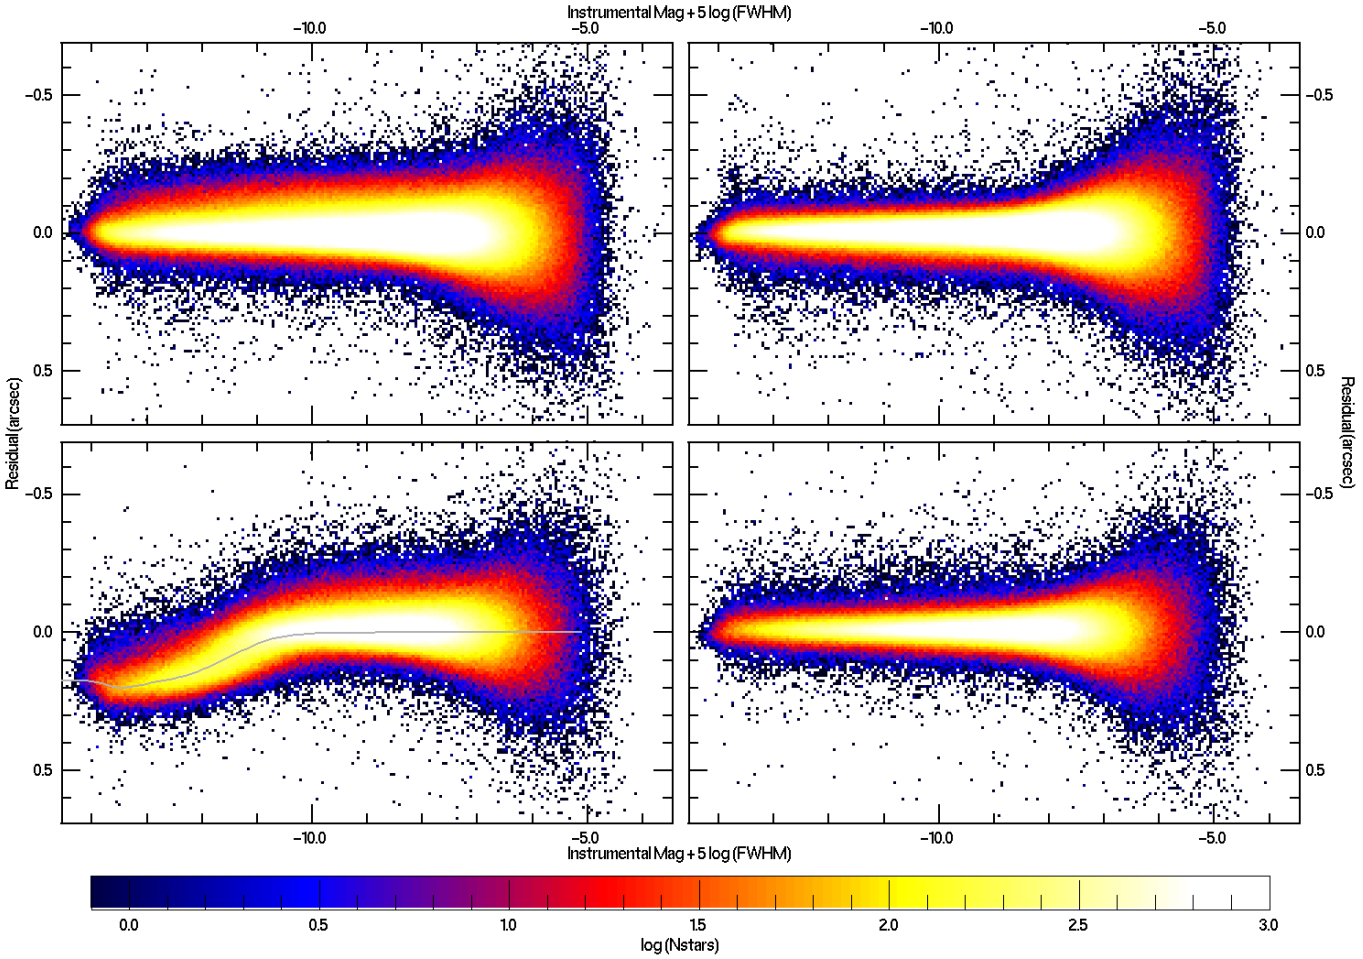


Figure 3. Illustration of the Koppenhöfer Effect on chip XY04. In each plot, the solid line shows the measured mean residual for stars detected on this chip as a function of the instrumental magnitude / FWHM^2. **top left** X-direction before correction. **top right** Y-direction before correction. **bottom left** X-direction after correction. **bottom right** Y-direction after correction.

the measured displacement in the direction perpendicular to the parallactic angle. We represent the trend with a spline fitted to this dataset.

The amplitude of the DCR trend in the five filters is $(g, r, i, z, y) = (0.010, 0.001, -0.003, -0.017, -0.021)$ arcsec airmass^{-1} magntiude^{-1}. We saturate the DCR correction if the term $colorTAN(\zeta)$ for a given measurement is outside a range where the DCR correction is well measured. The maximum DCR correction applied to the five filters is $(g, r, i, z, y) = (0.019, 0.002, 0.003, 0.006, 0.008)$ arcseconds.

7.1.3. *Astrometric Flat-field*

After correction for both KE and DCR, we observe persistent residual astrometric deviations which depend on the position in the camera. We construct an astrometric "flat-field" response by determining the mean residual displacement in the X and Y (chip) directions as a function of position in the focal plane. We have measured the astrometric flat using a sampling resolution of 40x40 pixels, matching the photometric flat-field correction images. Figures 6 and 7 show the astrometric flat-field images for the five filters $g_{\rm P1}$,$r_{\rm P1}$,$i_{\rm P1}$,$z_{\rm P1}$,$y_{\rm P1}$ in each of the two coordinate directions. These plots show several types of features.

The dominant pattern in the astrometric residual is roughly a series of concentric rings. The pattern is similar to the pattern of the focal surface residuals measured by (REF), which also has a concentric series of rings with similar spacing. The "tent" in the center of the focal surface reflected in these astrometry residual plots. Our interpretation of the structure is that the deviations of the focal plane from the ideal focal surface introduces small-scale PSF changes, presumably coupled to the optical aberrations, which result in small changes in the centroid of the object relative to the PSF model at that location. Since the PSF model shape parameters are only able to vary at the level of a 6x6 grid per chips, the finer structures are not included in the PSF model. The PV2 analysis shows the ring structure more clearly, with a pattern much more closely following the focal surface deviations. In the PV2 analysis, the PSF model used at most a 3x3 grid per chip to follow the shape variations, so any changes caused by the optical aberrations would be less well modeled in the PV2 analysis, as we observe.

A second pattern which is weakly seen in several chips consists of consistent displacements in the X (serial) direction for certain cells. This effect can be seen most clearly in chips XY45 and XY46. In the PV2 analysis, this pattern is also more clearly seen. In this case, the

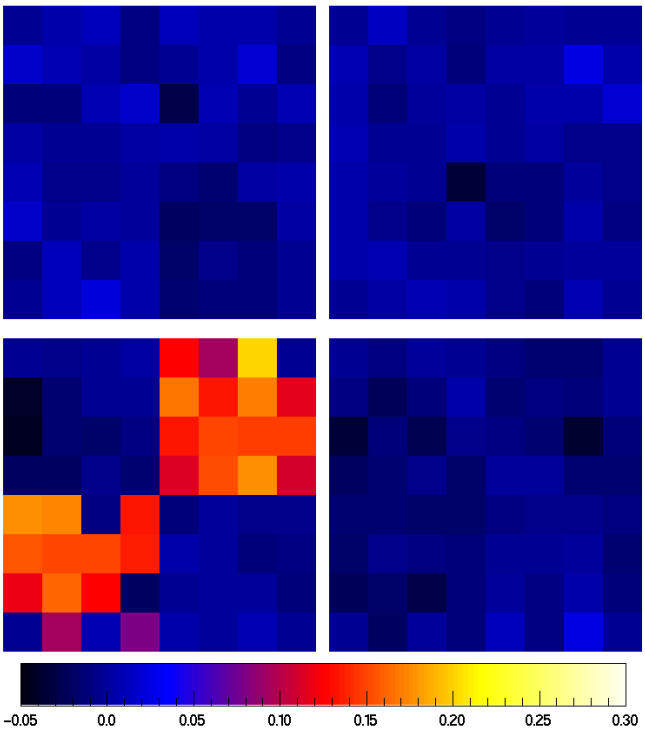


Figure 4. Map of the amplitude of the Koppenhöfer Effect on chips across the focal plane. In the affected chips, bright stars are up to 0.2 **arcsec** deviant from their expected positions. **bottom left** X-direction before correction. **bottom right** Y-direction before correction. **top left** X-direction after correction. **top right** Y-direction after correction.

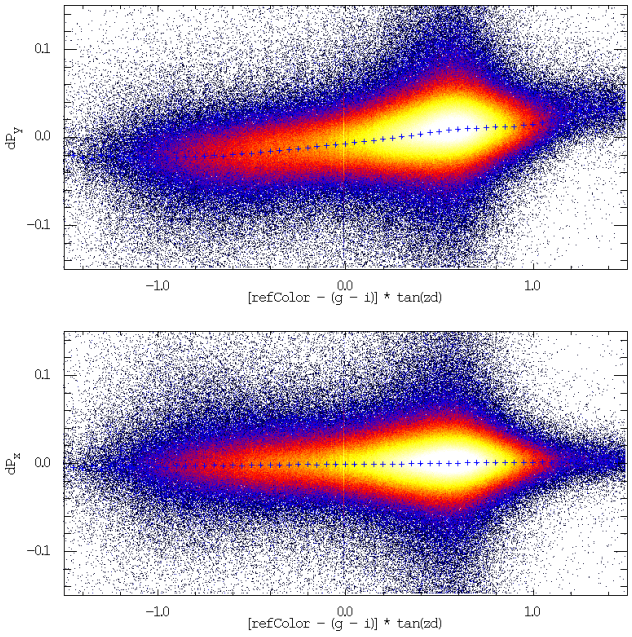


Figure 5. Example of the DCR trend in the g-band. **top:** DCR trend in the parallactic direction **bottom:** DCR trend perpendicular to the parallactic angle.

fact that the astrometric model used polynomials with a maximum of 3rd order per chip means the deviation of individual cells cannot be followed by the astrometric model.

A third effect is seen at the edge of the chips, where there appears to be a tendency for the residual to follow the chip edge. The origin of this is unclear, but likely caused by the astrometry model failing to follow the underlying variations because of the need to extrapolate to the edge pixels. Finally, we also mention an interesting effect *not* visible at the resolution of these astrometric flat-field images. Fine structures are observed at the ∼10 pixel scale similar to the "tree rings" reported by the DES team and others (G. Berstein REF & REFS). The details of these tree rings are beyond the scope of this article, and will be explored in future work.

Unfortunately, we discovered a problem with the astrometric flat-field correction too late to be repaired for DR1. As can be seen by inspection of Figures 6 and 7, there is significant pixel-to-pixel noise in the the astrometric flat-field images. This pixel-to-pixel noise is caused by too few stars used in the measuremnt of the flat-field structure for the high-resolution sampling. As a result, the astrometric flat-field correction reduces systematic structures on large spatial scales, but at the expense of degrading the quality of an individual measurement. Only i-band has sufficient signal-to-noise per pixel to avoid significantly increasing the per-measurement position errors.

Figure 8 shows the standard devitions of the mean residual astrometry in (α, δ) for bright stars as a function of position across the sky. For each pixel in these images, we selected all objects with $15 < i < 17$, with at least 3 measurements in i-band (to reject artifacts detected in a pair of exposures from the same night), with `PSF_QF` > 0.85 (to reject excessively-masked objects), and with $mag_{\rm PSF} - mag_{rmKron} < 0.1$ (to reject galaxies). We then generated histograms of the difference between the object position predicted for the epoch of each measurement (based on the proper motion and parallax fit) and the observed position of that measurement, in both the Right Ascension and Declination directions (in linear arcseconds), for all stars in a given pixel in the images. From these residual histograms, we can then determine the median and the 68%-ile range to calculate a robust standard deviation. This represents the bright-end systematic error floor for a measurement from a single exposure. The standard deviations are then plotted in Figure 2. The median value of the standard deviations across the sky is $(\sigma_\alpha, \sigma_\delta) = (22, 23)$ milliarcseconds.

The Galactic plane is clearly apparently in these images. Like photometry, we attribute this to failure of the PSF fitting due to crowding. The celestial North pole regions have somewhat elevated errors in both R.A. and DEC. This may be due to the larger typical seeing at these high airmass regions, but without further exploration this is interpretation uncertain. Several features can be seen which appear to be an effect of the tie to the Gaia astrometry: the stripes near the center of the DEC image and the right side of the R.A. image. The mesh of circular outlines is due to the outer edge of the focal plane where the astrometric calibration is poorly determined. As discussed above, the median values in the images are higher than expected based on our PV2 analysis of the

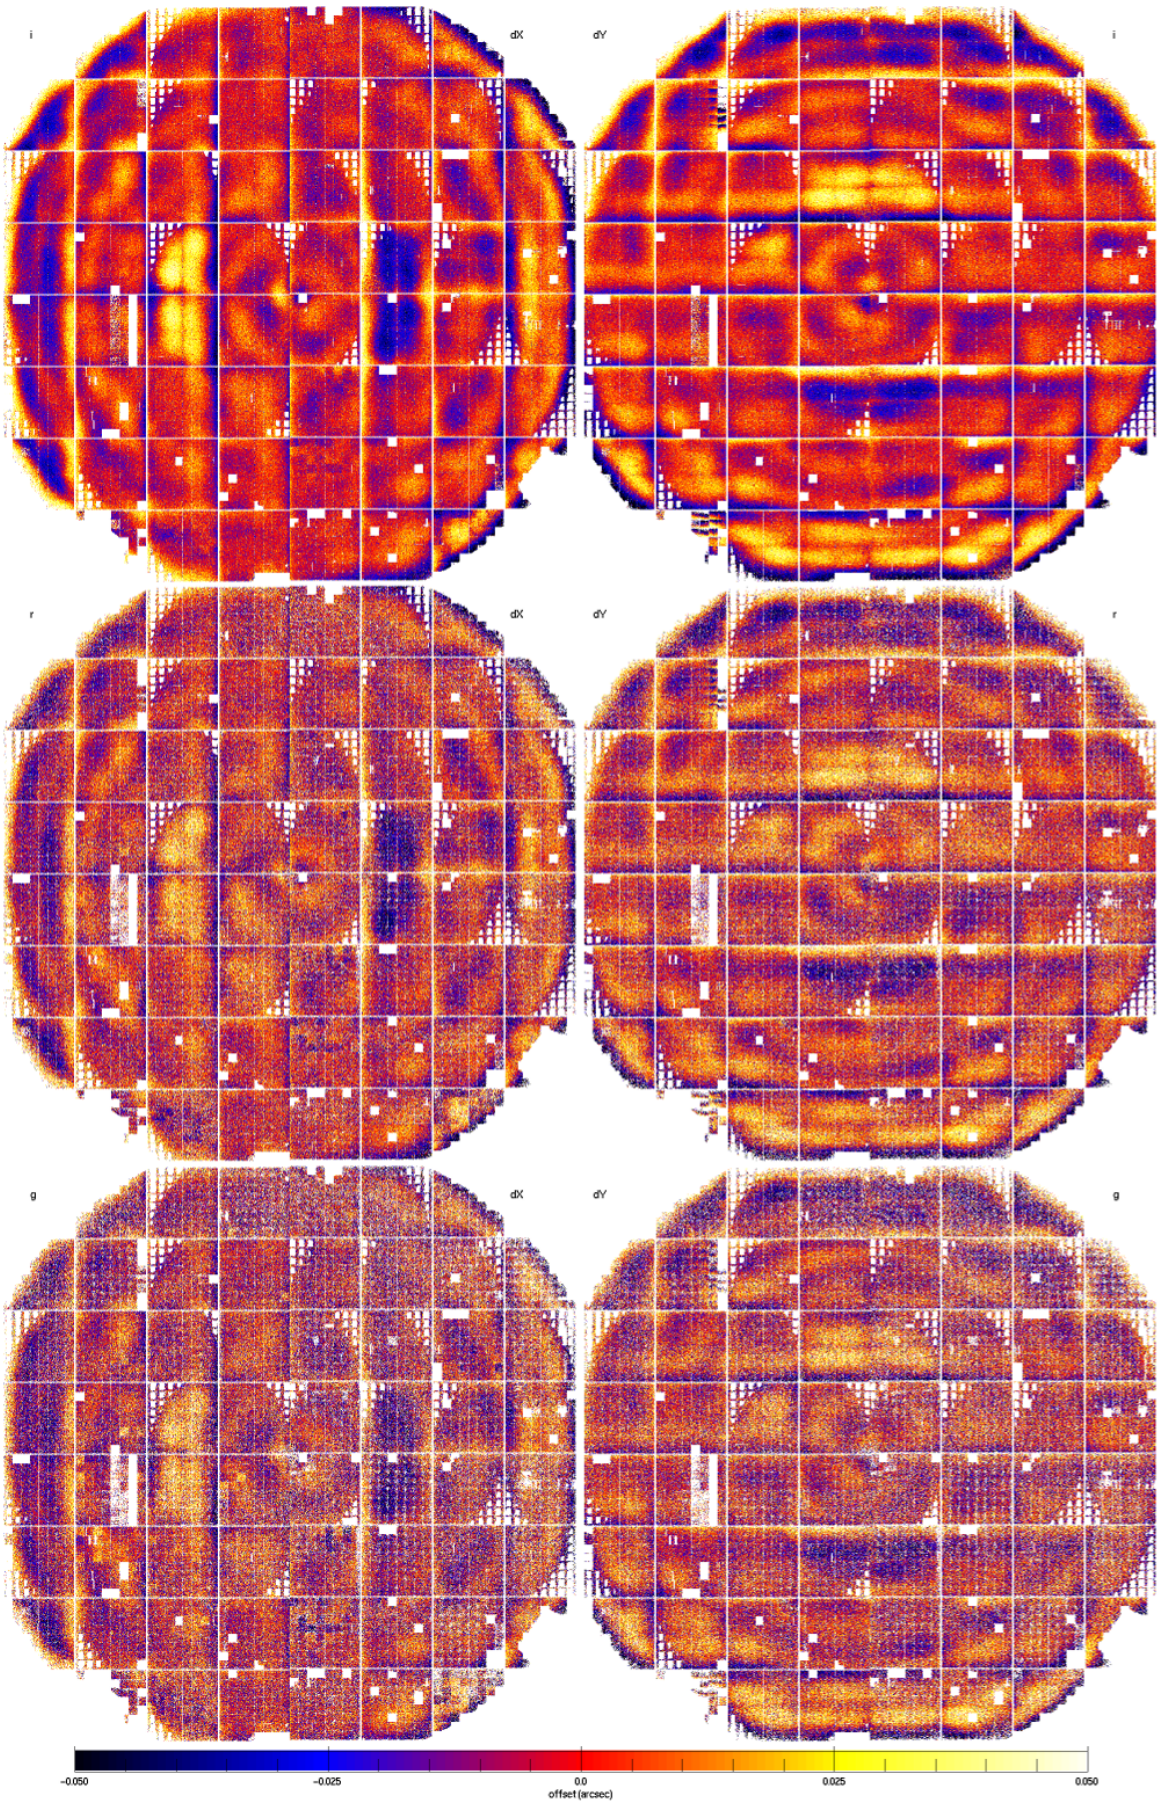


Figure 6. High-resolution astrometric flat-field correction images for *gri*.

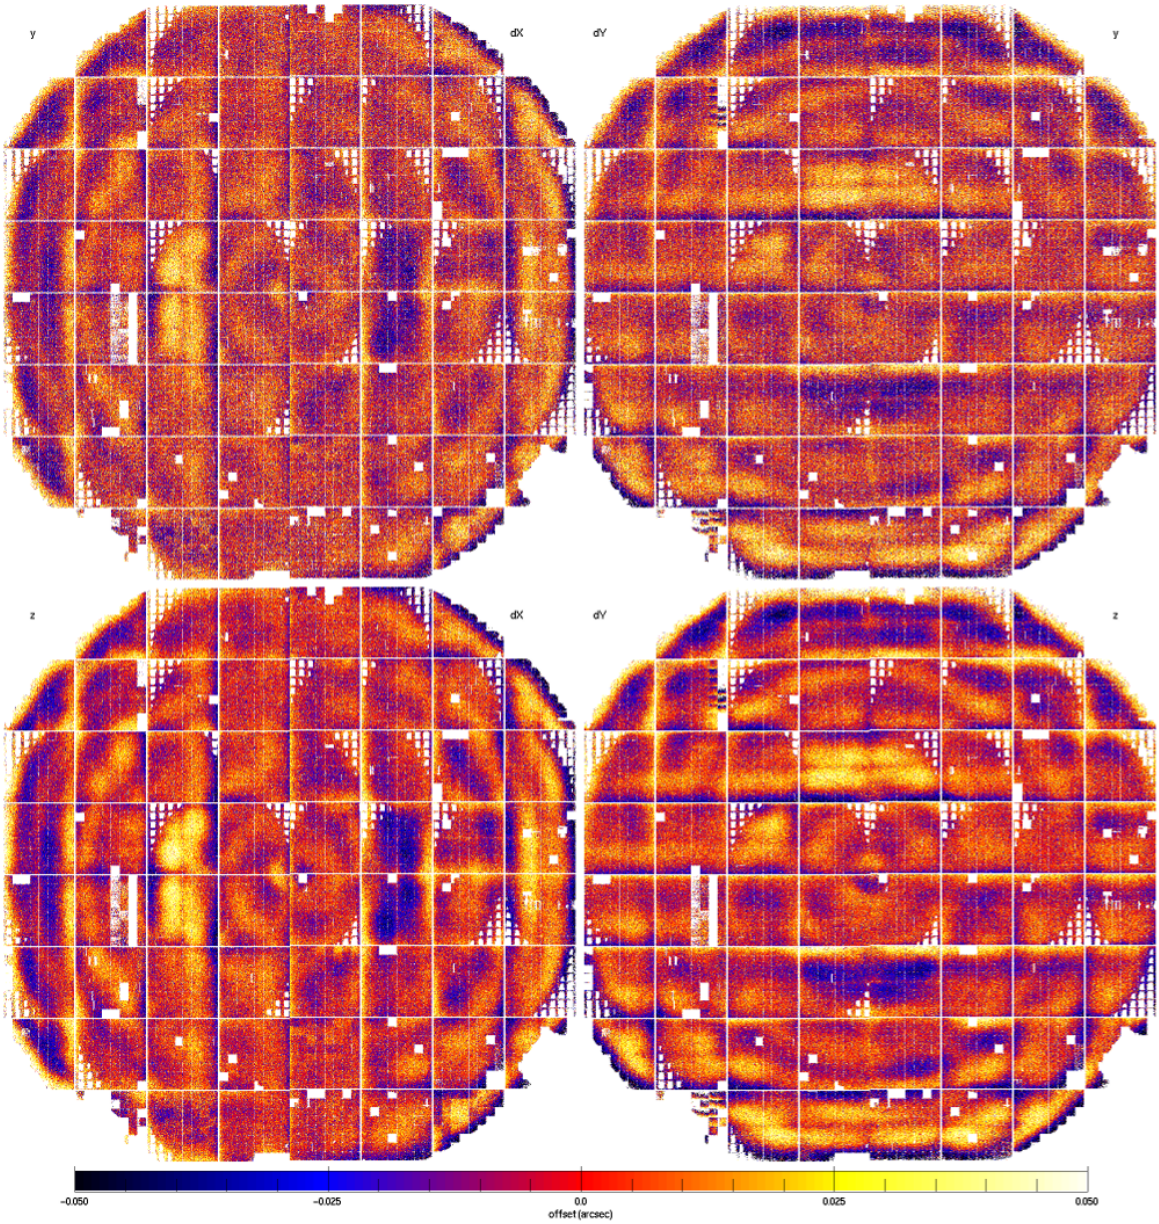


Figure 7. High-resolution astrometric flat-field correction images for zy.

astrometry: the median per-measurement error floor of ∼22 mas is significantly worse than the ∼17 mas value in that earlier analysis. We attribute this degradation to the noise introduced by the astrometric flat-field. This noise can likely be addressed before the DR2 release of the individual measurement data.

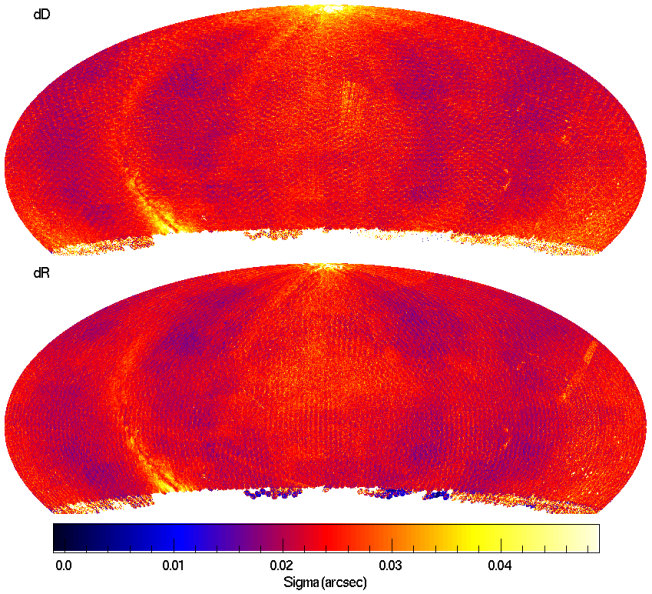


Figure 8. Consistency of photometry measurements across the sky. Each panel shows a map of the standard deviation of astrometry residuals for stars in each pixel. The median value of the standard deviations across the sky is $(\sigma_\alpha, \sigma_\delta) = (22, 23)$ milliarcseconds. These values reflect the typical single-measurement errors for bright stars. See discussion regarding the astrometric flat which is likely responsible for these elevated value.

After the initial analysis to measure the KE corrections, DCR corrections, and astrometric flat-field corrections, we applied these corrections to the entire database. Within the schema of the database, each measurement has the raw chip coordinates (`Measure.Xccd,Yccd`) as well as the offset for that object based on each of these three corrections: `Measure.XoffKH,YoffKH`, `Measure.XoffDCR,YoffDCR`, `Measure.XoffCAM,YoffCAM`. The offsets are calculated for each measurement based on the observed instrumental chip magnitudes and FWHM for the Koppenhoffer Effect, on the average chip colors and the altitude & azimuth of each measurement for the DCR correction, and on the chip coordinates for the astrometric flat-field corrections. The corrections are combined and applied to the raw chip coordinates and saved back in the database in the fields `Measure.Xfix,Yfix`. At this point, we are ready to run the full astrometric calibration.

7.2. *Galactic Rotation and Solar Motion*

The initial analysis of the PV2 astrometry used the 2MASS positions as an inertial constraint: the 2MASS coordiates were included in the calculation of the mean positions for the objects in the database, with weight corresponding to the reported astrometric errors. In this analysis, the object positions used to determine the calibrations of the image parameters ignored proper motion and parallax. After the image calibrations were determined, then individual objects were fitted for proper motion and possibly parallax, as discussed in detail below.

Using the PV2 analysis of the astrometry calibration, we discovered large-scale systematic trends in the reported proper motions of background quasars. This motion had an amplitude of 10 - 15 milliarcseconds per year and clear trends with Galactic longitude. We also observed systematic errors of the mean positions with respect to the ICRF milliarcsecond radio quasar positions, with an amplitude of ∼60 milliarcseconds, again with trends associated with Galactic longitude. Since the 2MASS data were believed to have minimal average deviations relative to the ICRF quasars, this latter seemed to be a real effect.

We realized that both the proper motion and the mean position biases could be caused by a single common effect: the proper motion of the stars used as reference stars between the 2MASS epoch (∼2000) and PS1 epoch (∼2012). Since we are fitting the image calibrations without fitting for the proper motions of the stars, we are in essencence forcing those stars to have proper motions of 0.0. The background quasars would then be observed to have proper motions corresponding to the proper motions of the reference stars, but in the opposite direction. We demonstrated that the observed quasar proper motions agreed well with the distribution expected if the median distance to our reference stars was ∼500 pc.

For PV3, we desired to address this bias by including our knowledge about the distances to the reference stars and the expected typical proper motions for stars at those distances. With some constraint on the distance to each star, we can determine the expected proper motion based on a model of the Galactic rotation and solar motions. We can then calculate the mean positions for the objects keeping the assumed proper motion fixed. When calibrating a specific image, the reference star mean position is then translated to the expected position at the epoch of that image. The image calibration is then performed relative to these predicted postions. This process naturally accounts for the proper motion of the reference stars. In order to make the calibrations consistent with the observed coordinates of an external inertial reference, we perform the iterative fits using the technique as described, but assign very high weights in the initial iterations to the inertial reference, and reduce the weights as the astrometric calibration iterations proceed.

In order to perform this analysis, we need estimated distances for every reference star used in the analysis. Green et al (REF) performed SED fitting for 800M stars in the 3π region using PV2 data. The goal of this work was to determine the 3D structure of the dust in the galaxy. By fitting model SEDs to stars meeting a basic data quality cut, they determined the best spectral type, and thus $T_{\rm eff}$, absolute r-band magnitude, distance modulus, and extinction A_V (the desired output and used to determine the dust extinction as a function of distance throughout the galaxy). We use the distance modulus determined in this analysis to predict the proper motions.

To convert the distances to proper motions, we use the Galactic rotation parameters $(A, B) = (14.82, -12.37)$ km sec^{-1} pc^{-1} and Solar motion parameters $(U_{\rm sol}, V_{\rm sol}, W_{\rm sol})$

= (9.32, 11.18, 7.61) km sec^{-1} as determined by Feast & Whitelock (REF) using Hipparchos data. Proper motions are determined from the following:

$$\mu_l^{\rm gal} = (A\cos(2l) + B)\cos(b) \quad (13)$$

$$\mu_b^{\rm gal} = \frac{-A\sin(2l)\sin(2b)}{2} \quad (14)$$

$$\mu_l^{\rm sol} = \frac{U\sin(l) - V\cos(l)}{d} \quad (15)$$

$$\mu_b^{\rm sol} = \frac{(U\cos(l) + V\sin(l))\sin(b) - W\cos(b)}{d} \quad (16)$$

where d is the distance and l, b are the Galactic coordintes of the star. Note that the proper motion induced by the Galactic rotation is independent of distance while the reflex motion induced by the solar motion decreases with increasing distance. Also note that this model assumes a flat rotation curve for objects in the thin disk; any reference stars which are part of the halo population will have proper motions which are not described by this model; the mostly random nature of the halo motions should act to increase the noise in the measurement, but should not introduce detectable motion biases. Also, if the distance modulus is not well determined, we can assume the object is simply following the Galactic rotation curve and set a fixed proper motion. If we do not have a distance modulus from the Green et al analysis, we assume a value of 500pc.

7.3. *Gaia Constraint*

After the full relative astrometry analysis was performed for the PV3 database, the Gaia Data Release 1 became available (**??**). This afforded us the opportunity to constrain the astrometry on the basis of the Gaia observations. Gaia DR1 objects which are bright enough to have proper motion and parallax solutions are in general saturated in the PS1 observations. Thus, we are limited to using the Gaia mean positions reported for the fainter stars. We extracted all Gaia sources not marked as a duplicate from the Gaia archive and generated a DVO database from this dataset. We then merged the Gaia DVO into the PV3 master DVO database. We reran the complete relative astrometry analysis using Gaia as an additional measurement. We applied the analysis described above, applying the estimated distances to determine preliminary proper motions. The Gaia mean epoch is reported as 2015.0, so all Gaia measurements were assigned this epoch. We wanted to ensure the Gaia measurements dominated the astrometric solutions, so we made the weight very high for the Gaia points: 1000× the nominal weight in the initial fits (to lock down the reference frame), decreasing to 100× the nominal weight for the last fits. We also retained the 2MASS measurements in the analysis, but gave them somewhat lower weights than Gaia: while the 2MASS data does not have the accuracy of Gaia, the coverage is known to be quite complete, while the Gaia DR1 has clear gaps and holes. Having 2MASS, even at a lower weight, helps to tile over those gaps.

Figure 9 shows a comparison between the Pan-STARRS photometry in g, r, i and the Gaia photometry in the G-band. To compare the PS1 photometry to the very broadband Gaia G filter, we have determined a transformation based on a 3rd order polynomial fit to $g - r$ and $g - i$ colors. This transformation reproduces Gaia photometry reasonably well for stars which are not too red. For a comparison, we have seleted all PS1 stars with Gaia measurements meeting the following criteria: $14 < i < 19$, with at least 10 total measurements, within a modest color range $0.2 < g-r < 0.9$. We also restricted to objects with $i_{\rm PSF} - i_{\rm Kron} < 0.1$, using the average i magnitudes determined from the individual exposures.

For Figure 9, we calculate the difference between the estimated G-band magnitude based on PS1 g, r, i photometry and the G-band photometry reported by Gaia. For each pixel, we determine the histogram of these differences and calculate the median and the 68%-ile range. In Figure 9, these values are plotted as a color scale.

The Galactic plane is clearly poorly matched between the two photometry systems. This may in part be due to the difficulty of predicting G-band magnitudes for stars which are significantly extincted: the G-band includes significant flux from the PS1 z-band which was not used in our transformation. Many other large scale feature in the median differences have structures similar to the Gaia scanning pattern (large arcs and long parallel lines. There are also structures related to the PS1 exposure footprint. These show up as a mottling on the ∼3 degree scale (e.g., lower right below the Galactic plane). The amplitude of the residual structures is fairly modest. The standard devition of the median difference values is 7 millimagnitudes. This number gives an indication of the overall photometric consistency of both Gaia and PS1 and implies that the systematic error floor for each survey is less than 7 millimags.

Figure 10 shows a comparison between the Pan-STARRS mean astrometry positions in α, δ and the Gaia astrometry. For this comparison, we have seleted all PS1 stars with Gaia measurements with $14 < i < 19$ and with at least 10 total measurements. For Figure 10, we calculate the difference between the position predicted by PS1 at the Gaia epoch (using the proper motion and parallax fit) and the position reported by Gaia. For each pixel, we determine the histogram of these differences in the R.Aând DEC directions, and calculate the median and the 68%-ile range. In Figure 10, these values are plotted as a color scale.

There is good consistency between the PS1 and Gaia astrometry. There are patterns from the Galactic plane (though not very strongly at the bulge). There are also clear features due to the PS1 exposure footprint (ring structure on ∼3 degree scales). In the plots of the scatter, there are patterns which are related to the Gaia scanning rule. These are presumably regions with relatively low signal to noise in Gaia; they were also apparent in the plots of the statisics of the per-exposure measurement residuals (Figure 8. The standard deviations of the median differences are $(\sigma_\alpha, \sigma_\delta) = (4, 3)$ milliarcseconds.

7.4. *Calculation of Object Astrometry*

7.4.1. *Iteratively Reweighted Least Squares Fitting*

7.4.2. *Seletion of Measurements*

8. DISCUSSION

9. CONCLUSION

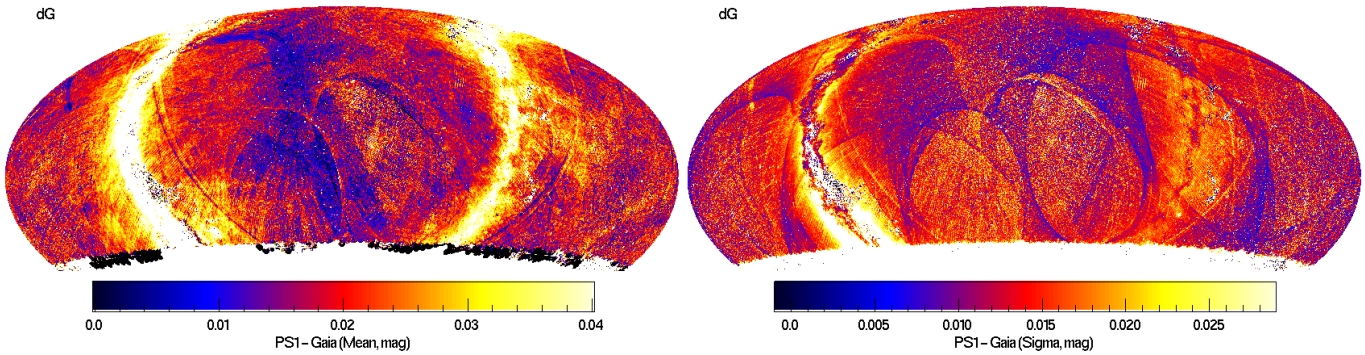


Figure 9. Comparison with Gaia photometry. **Left** Mean of PS1 - Gaia, **Right** Standard deviation of PS1 - Gaia. For pixels with $|b| > 30$ and $\delta > -30$, the standard deviation of the PS1 - Gaia mean values is 7 millimagnitudes, while the median of the standard deviations is 12 millimagnitudes. The former is a statement about the consistency of the Gaia and Pan-STARRS 1 photometry, while the latter reflects the combined bright-end errors for both systems.

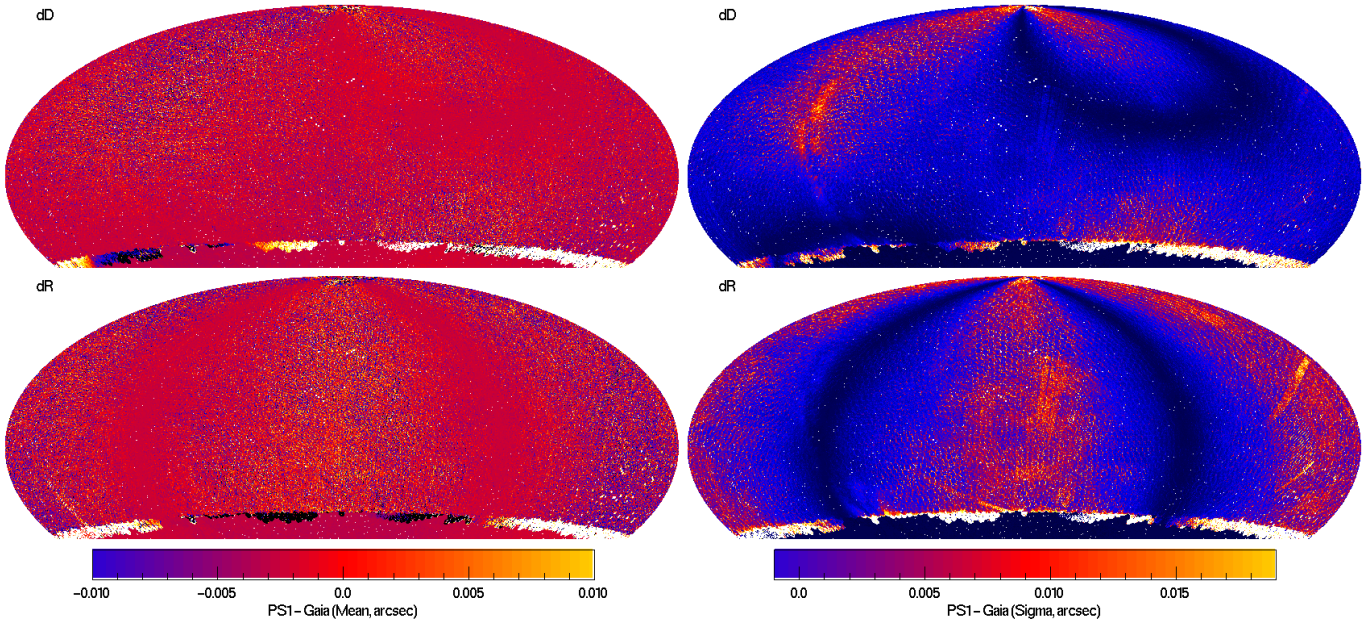


Figure 10. Comparison with Gaia astrometry. **Left** Mean of PS1 - Gaia, **Right** Standard deviation of PS1 - Gaia. The median value of the standard deviations is $(\sigma_\alpha, \sigma_\delta) = (4, 3)$ milliarcseconds.

The Pan-STARRS1 Surveys (PS1) have been made possible through contributions of the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation under Grant No. AST-1238877, the University of Maryland, and Eotvos Lorand University (ELTE) and the Los Alamos National Laboratory.

REFERENCES

Calabretta, M. R. & Greisen, E. W. 2002, A&A ,, 395, 1077
Chambers, K. C., Magnier, E. A., Metcalfe, N., & et al. 2017, ArXiv e-prints
Flewelling, H. A., Magnier, E. A., Chambers, K. C., Heasley, J. N., Holmberg, C., Huber, M. E., Sweeney, W., Waters, C. Z., Chen, T., Farrow, D., Hasinger, G., Henderson, R., Long, K. S., Metcalfe, N., Nieto-Santisteban, M. A., Norberg, P., Saglia, R. P., Szalay, A., Rest, A., Thakar, A. R., Tonry, J. L., Valenti, J., Werner, S., White, R., Denneau, L., Draper, P. W., Jedicke, R., Kudritzki, R.-P., Price, P. A., Chastel, S., McClean, B., Postman, M., & Shiao, B. 2016, ArXiv e-prints
Gaia Collaboration, Brown, A. G. A., Vallenari, A., Prusti, T., de Bruijne, J. H. J., Mignard, F., Drimmel, R., Babusiaux, C., Bailer-Jones, C. A. L., Bastian, U., & et al. 2016, A&A ,, 595, A2
Hodapp, K. W., Siegmund, W. A., Kaiser, N., Chambers, K. C., Laux, U., Morgan, J., & Mannery, E. 2004, in Proc. SPIE, Vol. 5489, Ground-based Telescopes, ed. J. M. Oschmann, Jr., 667–678
Huber, M., TBD, A., TBD, B., & et al. 2017, ArXiv e-prints

Juric, M. 2011, in Bulletin of the American Astronomical Society, Vol. 43, American Astronomical Society Meeting Abstracts #217, 433.19
Magnier, E. A., Schlafly, E., Finkbeiner, D., Juric, M., Tonry, J. L., Burgett, W. S., Chambers, K. C., Flewelling, H. A., Kaiser, N., Kudritzki, R.-P., Morgan, J. S., Price, P. A., Sweeney, W. E., & Stubbs, C. W. 2013, ApJS ,, 205, 20
Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., & et al. 2017, ArXiv e-prints
Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., Tonry, J. L., Goldman, B., Röser, S., Schilbach, E., Chambers, K. C., Flewelling, H. A., Huber, M. E., Price, P. A., Sweeney, W. E., Waters, C. Z., Denneau, L., Draper, P., Hodapp, K. W., Jedicke, R., Kudritzki, R.-P., Metcalfe, N., Stubbs, C. W., & Wainscoast, R. J. 2016a, ArXiv e-prints
Magnier, E. A., Sweeney, W. E., Chambers, K. C., Flewelling, H. A., Huber, M. E., Price, P. A., Waters, C. Z., Denneau, L., Draper, P., Jedicke, R., Hodapp, K. W., Kudritzki, R.-P., Metcalfe, N., Stubbs, C. W., & Wainscoast, R. J. 2016b, ArXiv e-prints
Onaka, P., Tonry, J. L., Isani, S., Lee, A., Uyeshiro, R., Rae, C., Robertson, L., & Ching, G. 2008, in Proc. SPIE, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, 70140D
Schlafly, E. F., Finkbeiner, D. P., Jurić, M., Magnier, E. A., Burgett, W. S., Chambers, K. C., Grav, T., Hodapp, K. W., Kaiser, N., Kudritzki, R.-P., Martin, N. F., Morgan, J. S., Price, P. A., Rix, H.-W., Stubbs, C. W., Tonry, J. L., & Wainscoat, R. J. 2012, ApJ ,, 756, 158
Scolnic, D., Casertano, S., Riess, A., Rest, A., Schlafly, E., Foley, R. J., Finkbeiner, D., Tang, C., Burgett, W. S., Chambers, K. C., Draper, P. W., Flewelling, H., Hodapp, K. W., Huber, M. E., Kaiser, N., Kudritzki, R. P., Magnier, E. A., Metcalfe, N., & Stubbs, C. W. 2015, ApJ ,, 815, 117
Scolnic, D., Rest, A., Riess, A., Huber, M. E., Foley, R. J., Brout, D., Chornock, R., Narayan, G., Tonry, J. L., Berger, E., Soderberg, A. M., Stubbs, C. W., Kirshner, R. P., Rodney, S., Smartt, S. J., Schlafly, E., Botticella, M. T., Challis, P., Czekala, I., Drout, M., Hudson, M. J., Kotak, R., Leibler, C., Lunnan, R., Marion, G. H., McCrum, M., Milisavljevic, D., Pastorello, A., Sanders, N. E., Smith, K., Stafford, E., Thilker, D., Valenti, S., Wood-Vasey, W. M., Zheng, Z., Burgett, W. S., Chambers, K. C., Denneau, L., Draper, P. W., Flewelling, H., Hodapp, K. W., Kaiser, N., Kudritzki, R.-P., Magnier, E. A., Metcalfe, N., Price, P. A., Sweeney, W., Wainscoat, R., & Waters, C. 2014, ApJ ,, 795, 45
Tonry, J. & Onaka, P. 2009, in Advanced Maui Optical and Space Surveillance Technologies Conference, E40
Tonry, J., Onaka, P., Luppino, G., & Isani, S. 2006, in The Advanced Maui Optical and Space Surveillance Technologies Conference, E47
Tonry, J. L., Burke, B. E., Isani, S., Onaka, P. M., & Cooper, M. J. 2008, in Proc. SPIE, Vol. 7021, High Energy, Optical, and Infrared Detectors for Astronomy III, 702105
Tonry, J. L., Stubbs, C. W., Lykke, K. R., Doherty, P., Shivvers, I. S., Burgett, W. S., Chambers, K. C., Hodapp, K. W., Kaiser, N., Kudritzki, R.-P., Magnier, E. A., Morgan, J. S., Price, P. A., & Wainscoat, R. J. 2012, ApJ ,, 750, 99
Waters, C. Z., Magnier, E. A., Price, P. A., Chambers, K. C., Draper, P., Flewelling, H. A., Hodapp, K. W., Huber, M. E., Jedicke, R., Kaiser, N., Kudritzki, R.-P., Lupton, R. H., Metcalfe, N., Rest, A., Sweeney, W. E., Tonry, J. L., Wainscoat, R. J., Wood-Vasey, W. M., & Builders, P. 2016, ArXiv e-prints